A REVEALED PREFERENCE THEORY FOR EXPECTED UTILITY

Edward J. Green and Kent Osband

January 1989

P-7485

The RAND Corporation, 1700 Main Street, P.O. Box 2138, Santa Monica, CA 90406-2138

A REVEALED PREFERENCE THEORY FOR EXPECTED UTILITY

by

Edward J. Green

Department of Economics
University of Pittsburgh

and

Kent Osband

RAND Corporation

ABSTRACT

Standard axiomatizations of expected-utility theory envision an agent with fixed probability assessments, who can be observed to choose actions from varying opportunity sets (for instance, pairs of lotteries). These axiomatizations also envision that the agent's preferences among these actions depend on the state of nature only through the state-dependent consequences of the actions, and that these consequences are clearly defined and observable. We suggest that this conception may be an unnecessarily restrictive basis for empirical testing, and instead study the pattern of choices from a fixed set of actions as probability assessments change. We show that maximization of the expectation of a general, state-dependent utility function places nontrivial restrictions on such a choice pattern. These restrictions are completely characterized by a discrete version of an integrability condition.

1. Introduction

There has been long-standing interest in testing expected-utility theory against alternative theories of choice in risky or uncertain situations. Until recently, the only direct tests were based on experimentally generated data. The results of many of these laboratory experiments were unfavorable to the expected-utility hypothesis.[1] Within the past decade, though, a substantial body of non-experimental evidence has been shown to conform well to the expected-utility hypothesis. This data concerns a variety of dynamic, stochastic choice problems related to optimal stopping. (These studies include Miller (1984), Pakes (1986), Rust (1987) and Wolpin (1984).)

These conflicting results could be reconciled in several different ways. It is possible that optimal-stopping problems are somehow psychologically special, and that people's performance with respect to these problems corresponds more closely to expected-utility maximization than does their general performance. It is also possible, though, that the experimental evidence systematically under-represents people's conformity to expected-utility theory. This under-representation may be especially severe with respect to decision making by experts. Expected-utility maximization with respect to an area of substantive expertise is a learned, non-verbal, domain-specific

[1] A particularly careful experiment of this sort was done by Grether and Plott (1979). References to other experiments are given in the bibliography of Allais (1979).

skill.[2] Such skill may have to be acquired slowly and through intensive study (for instance, over a period of several years of apprenticeship or postgraduate professional education), and it may not transfer instantaneously to an artificial task in an experiment, even after it has been acquired. These considerations show the great importance of further examining the expected-utility hypothesis outside of the laboratory, and particularly with respect to a variety of decision situations faced by experts.[3] In this paper, we provide a characterization of expected-utility maximization that may be useful for such investigations.

This characterization differs from the usual ones (such as Savage (1954)) in two respects that we believe will facilitate its empirical application. First, while the usual characterizations assume a fixed probability distribution of the

---

[2]The non-verbal aspect of expected-utility maximization was particularly emphasized by Savage (1954). If experimental protocols require subjects to rely to a large extent on their linguistic capabilities, then the experiments may not be representative of other decision-making problems. An observation bearing on this point has been made by Shortliffe and Buchanan (1984, p. 236), who note the difficulty of obtaining verbal probability estimates of diagnoses from experienced physicians. This observation need not entail that physicians' treatment decisions are inconsistent with expected-utility maximization. The physicians may be skilled at making treatment decisions, but yet have little facility at verbalizing the grounds for those decisions.

[3]The success of applied economic theories that incorporate the expected-utility hypothesis has sometimes been claimed to corroborate the hypothesis. However, Machina (1982) and Chew, Epstein and Zilcha (1986) show that comparative-statics predictions cannot discriminate between expected utility and a class of alternative decision criteria.

possible states of nature and allow the set of feasible actions to vary (typically, over pairs of alternative lotteries), here the set of actions will remain fixed and the probability measure will be allowed to vary. Second, while the usual characterizations stipulate that the utility of taking an action in a given state of nature depends only on a specified consequence of the action in that situation, here utility will be defined directly in terms of the action and the state of nature.[4]

To see why the probability measure needs to be allowed to vary, consider the problem of deciding whether an investor's demand for financial assets is consistent with expected-utility maximization. If the joint distribution of asset returns is fixed, then this question can be decided. (Cf. Green, Lau and Polemarchakis (1978), Green and Srivastava (1986) and Varian (1983).) Changes in asset demand would be imputed to variation of the investor's budget set as asset prices change. However, actual changes in asset prices are thought to be largely a consequence of changes in investors' perceptions of their joint return distributions. Therefore, demand theory based on the usual characterization of expected utility cannot be applied straigtforwardly.[5] In contrast, the empirical work on dynamic

[4] In order to avoid reference to consequences, the utility function must be made state contingent.State-contingent expected utility has previously been axiomatized, in a fromework that otherwise resembles Savage (1954), by Karni et al. (1983).

[5] Attanasio and Weber (1987) and Epstein and Zin (1987) have recently considered the precise role of expected-utility maximization in the determination of equilibrium asset prices under a rational-expectations assumption.

expected-utility maximization has been successful precisely because Bayes' theorem makes it possible to estimate how the agent's probability assessments evolve.

To see why avoidance of reference to consequences of actions may facilitate empirical work, consider an example of decision making in medicine. A state of nature specifies the situation of the patient, including the actual disease, state of general health, economic and family circumstances, and so forth. An action is a medical or surgical treatment. A consequence specifies all of the aspects of the doctor-patient relationship and of the patient's prognosis about which the doctor might be concerned, and that a treatment might affect. Given a patient's symptoms and medical history, the range of treatments that might reasonably be prescribed will usually be quite small. However, the consequence of a treatment is likely to be complex and difficult to measure precisely: is the progress of the disease arrested, is the patient well enough to resume normal activities, does the patient suffer pain or undesirable side effects, how high is the cost of treatment to the patient, how high is the doctor's remuneration, does the accomplishment of a state-of-the-art surgical procedure enhance the doctor's prestige in the medical community, and so forth. In this situation where observations about a narrow range of observed choices among acts would have to determine the utilities of a wide range of consequences, some aspects of which cannot be observed, the prospects for relating data about the doctor's decisions unambiguously to a utility function defined on consequences are poor.

This example of medical decisions provides a good example of how the present theory might be applied. Imagine a doctor who must base each treatment decision on evidence from a medical history and an examination of the patient. Suppose that, after the treatment is done, a completely accurate diagnosis becomes available. If the doctor treats many patients, then the conditional relative frequencies of these diagnoses should associate a probability measure on states of nature to each medical history and examination outcome. For each of these probability measures, the doctor has a preferred treatment. That is, the doctor's choices can be represented by an empirical behavior pattern that assigns actions to probability measures over states of nature.

Behavior patterns of this form are analyzed in this paper. It is shown that the hypothesis of expected-utility maximization places theoretical restrictions on the form of the behavior pattern. Whether or not a behavior pattern satisfies these restrictions can be determined by solving a system of linear inequalities. If the behavior is consistent with the expected-utility hypothesis, then the system has a solution that can be used to construct a utility function that rationalizes the behavior. These results bear some formal similarity to the standard problem of revealed preference in demand theory, and they could in fact be formulated within a generalized version of that theory such as Richter (1979) has provided.[6]

2. Choice of actions under uncertainty

The model to be studied here is a familiar one in most respects. There is a set of acts, and a set of states of nature. For mathematical simplicity, it will be assumed that both of these sets are finite. Let A be the acts, and $S=\{s^0,\ldots,s^n\}$ be the states. It is assumed that the subjective probability assessments of the subject may be determined by any probability measure over S corresponding to an element of a convex set $M\subseteq\Re^n$ having nonempty interior. If $\mu\in M$, then the probability measure corresponding to μ assigns probability μ_j to s^j for $j\geq 1$, and probability $1-\Sigma_{j=1}^{n}\mu_j$ to s^0.[7] It will be convenient to define $\zeta:M\to[0,1]$ by $\zeta(\mu)=1-\Sigma_{j=1}^{n}\mu_j$. Then, if $f:S\to\Re$, the expectation of f with respect to μ can be defined by

(1) $E_\mu[f]=\zeta(\mu)f(s^0)+\Sigma_{j\geq 1}\mu_j f(s^j)$.

A behavior pattern is a mapping $\beta:M\to A$. A utility function is a mapping $u:A\times S\to\Re$. For $a\in A$, define $u^a:S\to\Re$ by $u^a(s)=u(a,s)$. Utility function u rationalizes behavior pattern β if

(2) $\forall\mu\in M\ E_\mu[u^{\beta(\mu)}]=\max_{a\in A}E_\mu[u^a]$, and

---

[6] Another link between expected-utility theory and revealed-preference theory is provided by Border (1987), who investigates choice among lotteries. He shows that choices can be rationalized by expected utility if they respect stochastic dominance, provided that the set of available lotteries is closed under mixtures. Because the utilities of actions in various states are not assumed in this paper to be determined by monetary payoffs from the actions, these actions cannot be ranked by stochastic dominance.

[7] Clearly M must be a subset of $\{\mu|\ \forall j\ \mu_j\geq 0$ and $\Sigma_{j=1}^{n}\mu_j\ 1\}$.

(3) $\forall a \in \beta(M)\ \exists \mu \in M\ (\beta(\mu)=a \text{ and } \forall a' \neq a\ E_\mu[u^{a'}]<E_\mu[u^a])$.[8]

Note that, if ϕ is a permutation of A, and if $\beta'=\phi\cdot\beta$ and u rationalizes β, then there is a utility function u' that rationalizes β': $u'(\phi(a),s)=u(a,s)$. Therefore, whether or not β is rationalizable depends only on the partition $\{\beta^{-1}(a)|a\in A\}$. Thus (2) and (3) may be restated in terms of partitions. For Π a finite partition of M, define an indirect utility function for Π to be any function $v:\Pi\times S\rightarrow\Re$. In particular, if β is a behavior pattern, u is a utility function, and $\Pi=\{\beta^{-1}(a)|a\in A\}$, then $v(\beta^{-1}(a)),s)=u(a,s)$ is an indirect utility function for Π. With these definitions of Π and v, and defining $v^\pi(s)=v(\pi,s)$ for $\pi\in\Pi$, (2) and (3) translate respectively as

(4) $\forall\pi\in\Pi\ \forall\mu\in\pi\ E_\mu[v^\pi]=\max_{\pi'\in\Pi}E_\mu[v^{\pi'}]$, and

(5) $\forall\pi\in\Pi\ \exists\mu\in\pi\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]$.

When Π is any finite partition of M and conditions (4) and (5) are satisfied by Π and v, it will be said that v rationalizes Π and that Π is expected-utility consistent.

If Π were the partition of M induced by the expectation of some random variable X, then (4) would correspond to the definition in statistical decision theory of v to be a proper scoring rule for E[X]. Although the goal of the present research (to characterize the partitions for which (4) can be satisfied) is different from the goal of research on scoring rules (to

---

[8] Essentially, (3) states that the subject could not be assured to do as well if any action were deleted from A. Without (3), a constant utility function would rationalize every behavior pattern.

characterize the indirect utility functions for which (4) can be satisfied with respect to a given partition), the two studies are very closely related. In particular, Savage (1971) observed the fundamental role of convexity that will be exploited below. In the next section, Lemma 3 resembles the principal result of Osband and Reichelstein (1985). Osband (1986) has pointed out that (4) and (5) imply condition (9) of this paper, the convexity of $cl(\pi)$.

3. A necessary and sufficient condition

The main question to be answered in this paper is: when can a finite partition of M be rationalized by an indirect utility function? A preliminary answer is given in this section. This answer will lead to a necessary and sufficient condition that can be formulated in terms of simultaneous satisfiability, in the strictly positive orthant of a Euclidean space, of a finite set of linear equations. This algebraic formulation plausibly will lead to the formulation of an efficient algorithm to decide the question.

The idea of the theorem is given by the following heuristic argument. Suppose that v is an indirect utility function for Π, a finite partition of M. Define $V:M\rightarrow\mathbb{R}$ by

(6) $V(\mu)=\max_{\pi'\in\Pi}E_{\mu}[v^{\pi'}]$.

Since $E_{\mu}[v^{\pi'}]$ is an affine function of μ for each π', by (1), V is a polyhedral convex function. (That is, V is both convex and piecewise linear. It is convex because it is the pointwise supremum of a set of affine functions, and piecewise linear because the set of affine functions is finite.) If (5) holds, then, by the continuity of $E_\mu[v^{\pi'}]$ for each $\pi'$and by the finiteness of Π, $\forall\pi\epsilon\Pi$ $\mathrm{int}(\pi)\neq\emptyset$.[9] Let $\mu\epsilon\mathrm{int}(\pi)$. Ther if (4) holds, the gradient of V at μ is

(7) $\nabla V(\mu)=(v(\pi,s^1)-v(\pi,s^0),\ldots,v(\pi,s^n)-v(\pi,s^0))$.[10]

This suggests that the interiors of the elements of Π are the largest sets on which the gradient of V is defined and constant. Moreover, given any polyhedral function $V:M\rightarrow\Re$ having a gradient that is related to Π in this way, (7) should determine an indirect utility function v that satisfies (6).

This conjecture is true, and is stated below as a theorem. To facilitate the statement of the theorem and later results, the conditions that have just been described will be characterized more formally. Suppose that

(8) Δ is a finite cover of M by polyhedral convex bodies in M.[11]

Furthermore, suppose that Δ is virtually a partition of M, except that the boundaries of its elements may overlap. That is, suppose that

(9) $\forall\delta\epsilon\Delta\ \forall\delta'\epsilon\Delta$ [If $\delta\neq\delta'$, then $\mathrm{int}(\delta)\cap\mathrm{int}(\delta")=\emptyset$.]

---

[9]Topological interior, closure, and boundary are taken relative to M.

[10]∇V denotes the gradient of V. Cf. Rockafellar (1970), p. 241.

[11]That is, (a) Δ is a finite collection of subsets of M, (b) M is the union of this collection, (c) each of the subsets is the intersection of M with a finite set of closed half-spaces, and (d) each subset has nonempty interior in M.

If (8) and (9) are satisfied, then call Δ a polyhedral decomposition of M. From the discussion above, it should be clear that a polyhedral convex function induces a polyhedral decomposition of its domain. Specifically, the function V induces the polyhedral decomposition Δ defined by

(10) $\delta\epsilon\Delta$ iff $\exists x\epsilon\mathbb{R}^n$ $[\delta = cl(\{\mu|\nabla V(\mu)=x\})]$.[12]

A polyhedral decomposition that satisfies (10) for some polyhedral convex function V will be called a subgradient decomposition. Now, the conjecture stated above can be reformulated as the assertion that Π is expected-utility consistent if and only if $\{cl(\pi)|\pi\epsilon\Pi\}$ is a subgradient decomposition. This assertion is now proved via several lemmas.

<u>Lemma 1</u>: If v rationalizes Π, then

$\pi\subseteq cl\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}$.

Proof: Let $\mu'\epsilon\pi$, and let μ satisfy $\forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]$. μ exists by (5). Define $\mu''=\alpha\mu+(1-\alpha)\mu'$, where $\alpha\epsilon(0,1)$. Then μ'' satisfies $\forall\pi'\neq\pi\ E_{\mu''}[v^{\pi'}]<E_{\mu''}[v^\pi]$. The lemma follows by letting α approach 0. Q.E.D.

<u>Lemma 2</u>: If v rationalizes Π, then

$int(\pi)=\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}$.

Proof: $\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}$ is open because Π is finite and because $E_\mu[v^{\pi'}]$ is continuous in μ for every π', so by

---

[12]Condition (10) means that $int(\pi)$ is the subset of M on which $\nabla V(\mu)$ is the unique subgradient of V. Recall that y is a subgradient of the convex function f at x if, for all x', $f(x')\geq f(x)+y(x'-x)$. This inequality is called Fenchel's inequality. It holds with strict inequality if y is not a subgradient of f at both x and x'.

Lemma 1, $\text{int}(\pi)\subseteq\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}$. The inclusion $\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}\subseteq\text{int}(\pi)$ follows from (4) because $\{\mu|\ \forall\pi'\neq\pi\ E_\mu[v^{\pi'}]<E_\mu[v^\pi]\}$ is open. Q.E.D.

<u>Lemma 3</u>: If v rationalizes Π, if V is defined by (6), and if $\mu\in\text{int}(\pi)$, then $\nabla V(\mu)$ is defined and satisfies (7). Suppose also that $\mu'\in\text{int}(\pi')$. Then

(11) $E_\mu[v^{\pi'}]=v(\pi',s^0)+\nabla V(\mu')\mu$.

Proof: By (4) and (6), V coincides with an affine function on a neighborhood of μ. Thus $\nabla V(\mu)$ is defined, and (7) is derived by computation using (1). The second assertion follows directly from (4), (6), and (7). Q.E.D.

<u>Lemma 4</u>: Suppose that v rationalizes Π, and that V is defined by (6). If $\pi\neq\pi'$, $\mu\in\text{int}(\pi)$, and $\mu'\in\text{int}(\pi')$, then $\nabla V(\mu)\neq\nabla V(\mu')$.

Proof: By Lemma 2 and (11), $0<[v(\pi,s^0)-v(\pi',s^0))]+[\nabla V(\mu)-\nabla V(\mu')]\mu$. Interchanging μ and μ', and π and π', and adding the resulting equation to the original one, yields $0<[\nabla V(\mu)-\nabla V(\mu')](\mu-\mu')$. This requires that $[\nabla V(\mu)-\nabla V(\mu')]\neq 0$. Q.E.D.

<u>Lemma 5</u>: Suppose that v rationalizes Π and that V is defined by (6). Then V is convex. If $\mu\in\text{int}(\pi)$, then $\text{int}(\pi)=\{\mu'|\nabla V(\mu)$ is the unique subgradient of V at $\mu'\}$.

Proof: Convexity of V is immediate from (6). Furthermore, $\text{int}(\pi)\subseteq\{\mu'|\nabla V(\mu)$ is the unique subgradient of V at $\mu'\}$, by Lemma 3. Now suppose that $\mu''\notin\text{int}(\pi)$. By Lemmas 1 and 2, there is a sequence $\{\mu^t\}\to\mu''$, with $\{\mu^t\}\subseteq\cup_{\pi'\neq\pi}\text{int}(\pi')$. By the finiteness of Π there is some $\pi'\neq\pi$ such that an infinite subsequence of $\{\mu^t\}$ is in $\mathrm{int}(\pi')$. Let $\mu'\in\mathrm{int}(\pi')$. By (7), $\nabla V(\mu^t)\to\nabla V(\mu')$ along this subsequence. Therefore, both $\nabla V(\mu)$ and $\nabla V(\mu')$ are subgradients of V at μ'', by Theorem 24.4 of Rockafellar (1970). This shows that $\{\mu' \mid \nabla V(\mu)$ is the unique subgradient of V at $\mu'\}\subseteq\mathrm{int}(\pi)$. Q.E.D.

Lemma 6: Suppose that $V:M\to\Re$ is a polyhedral convex function, and that Π and V satisfy (8) and (10). Then there is an indirect utility function v that rationalizes Π.

Proof: Let $\pi\in\Pi$. By (8), $\mathrm{int}(\pi)$ is nonempty. Let $\mu\in\mathrm{int}(\pi)$. In order to satisfy (11) and (7), respectively, define $v(\pi,s^0)=V(\mu)-\nabla V(\mu)\mu$, and define $v(\pi,s^j)=[\nabla V(\mu)]_j+v(\pi,s^0)$. Then, if $\mu'\in\mathrm{int}(\pi')$, $E_\mu[v^{\pi'}]=\zeta(\mu)[V(\mu')-\nabla V(\mu')\mu']+\Sigma_{j\geq1}\mu_j\{[\nabla V(\mu')]_j+[V(\mu')-\nabla V(\mu')\mu']\}=V(\mu')+\nabla V(\mu')(\mu-\mu')$. Thus $E_\mu[v^\pi]=V(\mu)$, and $E_\mu[v^{\pi'}]<V(\mu)$ if $\pi'\neq\pi$ by Fenchel's inequality and (10).[13] This establishes (5). (4) is established by a limit argument, using (8) and the continuity of V. Q.E.D.

Theorem 1: Π is expected-utility consistent if and only if $\{\mathrm{cl}(\pi) \mid \pi\in\Pi\}$ is a subgradient decomposition.

The proof of Theorem 1 is immediate from the lemmas.

4. Some partitions that are not expected-utility consistent

The requirement that an expected-utility-consistent partition should essentially coincide with a polyhedral decomposition is intuitive and geometrically concrete. Convexity is the key economic aspect; it means that if the same act is

[13] Fenchel's inequality is the statement that, for a convex function f, $f(x')\geq f(x)+y(x'-x)$, with strict inequality if y is not a subgradient of f at x'.

preferred at both μ^1 and μ^2, then that act will be preferred for any random draw between μ^1 and μ^2.

The additional requirement (10) imposed on a subgradient decomposition is less intuitive and more abstract. This section attempts to foster some intuition for the role of (10) by presenting two polyhedral decompositions that fail to be subgradient decompositions, and by showing that these cannot be obtained from expected-utility-consistent partitions. Both of these examples assume three states of nature (n=2), allowing M to be graphed (by making an affine transformation) as a neighborhood of the origin in the plane.

<u>Example 1</u>: The first example is graphed in Figure 1. It consists of three polyhedral convex sets. δ^1 is the intersection of M with the lower half plane, and δ^2 and δ^3 are the intersections of M with the right and left quadrants of the upper half plane, respectively.

Suppose that these three sets were the images of the intersections of M with the parts of the probability simplex where acts 1, 2, and 3 were weakly preferred by an expected-utility-maximizing agent. If act 3 were stricken from the choice set, then δ^3 would have to be "divided up" between δ^1 and δ^2 so as to leave a polyhedral decomposition. Clearly, the only way to do this would be to extend δ^2 to cover the entire upper half of the plane. It follows that the agent would be indifferent between act 1 and act 2 along the same border that was originally the line of indifference between act 2 and act 3. By transitivity, the agent would have to be indifferent between act 1 and act 2 along that border. But this cannot be, since that border is not contained in δ^1. Therefore the decomposition is not induced by an expected-utility-consistent partition.

Example 2: The second example is shown in Figure 2. It is constructed from two concentric squares. Let the vertexes of the smaller square have coordinates $(\pm 1, \pm 1)$, and let the vertexes of the larger square have coordinates $(\pm 3, \pm 3)$. Connect (1,1) to (3,2), (-1,1) to (-2,3), (-1,-1) to (-3,-2), and (1,-1) to (2,-3). This construction divides the larger square into five polyhedral convex sets. Let the set to the right of the smaller square be δ^0, and number the other sets outside the smaller square in counterclockwise order. Let the smaller square itself be δ^4.

This polyhedral decomposition can be shown not to be induced by any expected-utility-consistent partition by a similar method as above. If act 4 were eliminated from the choice set and if the agent were an expected utility maximizer, then δ^4 would have to be divided up among the other four regions so as to leave a polyhedral decomposition. But in fact even enlarging δ^0 through δ^4 as much as convexity will allow leaves a square "hole" around the origin, as Figure 3 shows. Hence the decomposition cannot be induced by an expected-utility-consistent partition. In economic terms, there is a circularity of preferences in the interior of the hole: Act 0 is strictly preferred to act 1, act 1 is strictly preferred to act 2, act 2 is strictly preferred to act 3, and act 3 is strictly preferred to act 0.

Incidentally, this example shows that being a subgradient decomposition is not always stable under small perturbations. Consider the 5-element decomposition formed by connecting each vertex of the smaller square to the corresponding vertex of the larger square. This is the subgradient decomposition induced by the function V defined by $V(x,y)=0$ if $\max(|x|,|y|)\ 1$, and $V(x,y)=\max(|x|,|y|)-1$ otherwise. Clearly, though, an arbitrarily small rotation of each of the diagonal edges of this decomposition will transform it to one that is qualitatively like the decomposition of the example. It is conceivable that this sort of instability might lead to inappropriate rejection of the hypothesis of expected-utility maximization due to measurement error.

5. Characterization of subgradient decompositions

The foregoing examples provide some intuition for why only subgradient decompositions, rather than all polyhedral decompositions, can arise from expected-utility maximization. However, the examples do not directly provide a concrete picture of what subgradient decompositions look like or a computational procedure for recognizing them. The problem is that condition (10) requires the existence of an unspecified polyhedral convex function. For any given polyhedral decomposition, we would like to be able to recognize whether a function satisfying the condition exists, and to be able to construct it if it does exist. Further analysis of the examples in this section will lead to such a constructive characterization.

Two observations are in order. First, both of the examples have been studied in the same way: by removing elements of the polyhedral decomposition, trying to enlarge the remaining elements to fill the gaps without violating the convexity restrictions imposed by expected-utility maximization, and finally checking for violations of transitivity. It is plausible that this strategy could be generalized, and that it would provide a characterization of the sort that we are seeking. However, we will derive a somewhat different characterization below.

Second, in the examples the impossibility of satisfying (10) has been quite apparent, but in general this issue will be more subtle. A third example, closely related to Example 1, will make this point clear. To motivate this example, consider how it can be shown by contradiction that Example 1 does not satisfy (10). Imagine that (10) held, and think of Figure 1 as a view from above of the polyhedral convex function V. To satisfy (10) all of the flats of this graph must have different slopes. The flats lying above δ^2 and δ^3 are obtained by creasing the graph above the boundary of $\delta 1$, so that the flats lying above δ^2 and δ^3 become "flaps" attached to the flat lying above δ^1. In order for those flaps to have different slopes, though, the surface of the graph must be "torn" above the boundary between δ^2 and δ^3. This cannot happen, because it would imply that V were discontinuous.

Now, consider a more subtle example for which continuity considerations do not suffice to show the impossibility of satisfying (10).

<u>Example 3</u>: This example is graphed in Figure 4. In place of the boundary between the two subsets of the upper half plane in Example 1, there is now an additional, wedge-shaped subset. The graph of V can be adjusted above this wedge to repair the discontinuity in the earlier example. However, the resulting function will not be convex.

To present this example algebraically, define $\delta^1=\{(x,y)|y\ 0\}$, $\delta^2=\{(x,y)|0\ y\ x\}$, $\delta^3=\{(x,y)|0\ y\ -x\}$, and $\delta^4=\{(x,y)|y\geq\max[x,-x]\}$. These four sets comprise a polyhedral decomposition Δ. The interiors of the elements of Δ are the regions on which the gradients of a piecewise-linear function V is defined and constant. Specifically, let $V(x,y)=0$ on δ^1, $V(x,y)=y$ on δ^2, $V(x,y)=-y$ on δ^3, and $V(x,y)=x$ on δ^4. However V is not convex, as is easily seen by considering $V(x,1)$ as a function of x.

Of course, the fact that a polyhedral decomposition Δ is induced by a non-convex function does not mean that it cannot also be induced by some other function that is convex.[14] In that case, Δ would satisfy (10) and be a subgradient decomposition. The same kind of argument as was given for Example 1 will work here also, but two of the elements of Δ -- either δ^2 and δ^4, or

---

[14] We were led to formulate Example 3 by J.-P. Benoit,who raised the question of whether (10) is more restrictive than it would be if only the piecewise linearity of V were required.

else δ^3 and δ^4, will have to be deleted before trying to extend the remaining subset of the upper half plane, if transitivity is to lead to a contradiction. Moreover it is obvious that, if δ^4 were to be subdivided into m wedges radiating from the origin, then all of these plus either δ^2 or δ^3 -- a total of m+1 elements of the polyhedral decomposition -- would have to be deleted before transitivity would lead to a contradiction. Thus the strategy of identifying violations of (10) by making deletions can become very complex, even when there are only three states of nature. It would be desirable to have a criterion that is of bounded complexity. Such a criterion, involving only the inspection of pairs of adjacent elements of Δ, is now provided.

Notice that in both Example 1 and Example 3, the nonnegative ray of the x-axis is a face of δ^2, but is not an entire face of δ^1. It turns out that such a mismatch of faces must not occur, if (10) is to hold. This no-mismatch condition can be reformulated in such a way that a graph can be associated with Δ if it is satisfied, and this graphical interpretation will be useful later.

To state the condition in this way, let δ be a polyhedral convex subset of $M \subseteq \mathbb{R}^n$ with nonempty interior. Define an edge of δ to be the relative interior of an (n-1)-dimensional face of δ. Define the polyhedral decomposition Δ of M to be a polyhedral graph if, whenever ω is an edge of δ and ω' is an edge of δ' and $\omega \cap \omega' \neq \emptyset$, then $\omega = \omega'$. Note that, according to these definitions, the elements of Δ are vertexes of a finite graph, and two such vertexes are adjacent in the graph if and only if they share an (n-1)-dimensional face in the decomposition.

<u>Theorem 2</u>: Every subgradient decomposition of M is a polyhedral graph.

Proof: Suppose that Δ is a subgradient decomposition. By (10), a function $c:\Delta\rightarrow\Re$ and a 1-1 function $g:\Delta\rightarrow\Re^n$ exist such that

(12) $\forall\mu\epsilon M\ V(\mu)=\max_{\delta\epsilon\Delta}[c(\delta)+g(\delta)\mu]$, and

(13) $\forall\delta\epsilon\Delta\ \delta=\{\mu|V(\mu)=c(\delta)+g(\delta)\mu\}$.

A contradiction will be derived from the supposition that Δ is not a polyhedral graph. Suppose particularly that ω is an edge of δ, ω' is an edge of δ', $\mu'\epsilon\omega\cap\omega'$, and $\omega\not\subseteq\omega'$. $\omega\cap\omega'$ is a relatively open set of a hyperplane H. Since (13) holds everywhere on this set with respect to both δ and δ', $c(\delta)=c(\delta')$ and $[g(\delta')-g(\delta)]$ is perpendicular to H. $\omega\backslash\omega'$ has nonempty relative interior in H, so there must be a third element δ'' of Δ having an edge ω'' and a member μ'' such that $\mu''\epsilon\omega\cap\omega''$. Again, it must be that $c(\delta)=c(\delta'')$ and $[g(\delta'')-g(\delta)]$ is perpendicular to H. That is, if p is the unique unit vector perpendicular to H and such that $p(\delta-H)\subseteq(-\infty,0]$, then $[g(\delta')-g(\delta)]$ and $[g(\delta'')-g(\delta)]$ are scalar multiples of p by positive constants α' and α'', respectively. Now, for some positive constant γ, $\mu'+\gamma p\epsilon \mathrm{int}(\delta')$ and $\mu''+\gamma p\epsilon \mathrm{int}(\delta'')$. For $\mu=\mu'+\gamma p$, (12) and (13) imply that $\alpha'\geq\alpha''$, and for $\mu=\mu''+\gamma p$, these equations imply that $\alpha''\geq\alpha'$. That is, $\alpha'=\alpha''$, contradicting g being 1-1. Q.E.D.

Theorem 2 states that being a polyhedral graph is a necessary condition for a polyhedral decomposition to be a subgradient decomposition, which is what is related to expected-utility maximization via Theorem 1. However, Example 2 shows that being a polyhedral graph is not sufficient to be a subgradient decomposition. Figure 2 shows clearly that Example 2 is a polyhedral graph, but it has already been shown not to satisfy (10). Further consideration of Example 2 will lead to the formulation of an additional necessary condition for a polyhedral graph to be a subgradient decomposition, and it will turn out that this new condition is sufficient as well as being necessary. The proof of sufficiency will be constructive. Specifically, it will show that the condition is equivalent to the consistency of a finite system of weak linear inequalities, and it will provide an algorithm for defining a polyhedral convex function V satisfying (10) if the condition holds.

We are now going to provide a different argument that Example 2 does not satisfy condition (10). This argument relies heavily on the fact that Example 2 is a polyhedral graph.[15] For

---

[15]The argument could be reformulated to apply to any polyhedral decomposition. To do so, associate a finite graph with the decomposition by stipulating that δ and δ' share an edge, the relative interior of $\delta \cap \delta'$, whenever that set has dimension n-1. Then the present argument and the proof of the characterization theorem (Theorem 3) below would remain sound without any essential change, and Theorem 2 could be derived from Theorem 3 as a corollary. That approach would have the advantage of making as clear as possible how fundamental are the integrability considerations with which Theorem 3 is concerned. We have proved Theorem 2 directly because it is of some independent interest to pursue as far as possible the elementary approach to characterization (i.e., without transforming the problem into one regarding consistency of a system of linear inequalities) that has been studied so far in the paper. In addition, the direct proof is actually shorter than is the proof from Theorem 3. Theorem 2 shows that there is no loss of

any distinct δ^j and δ^k that share an edge, let ϵ^{jk} be a nonzero vector of $\mathfrak{R}^2$ satisfying

(14) $\epsilon^{jk}(\delta^k-\delta^j)\subseteq[0,\infty)$.

ϵ^{jk} is called the directed edge from δ^j to δ^k. Let E be an assignment of these vectors, one vector to each ordered pair of vertexes that share an edge. (That is, if δ^j and δ^k are adjacent, then E will specify both non-zero vectors ϵ^{jk} and ϵ^{kj}.) For any polyhedral graph Δ and set E of vectors satisfying these conditions (where $E\subseteq\mathfrak{R}^n$ if $M\subseteq\mathfrak{R}^n$), call (Δ,E) a polyhedral directed graph. If δ^j and δ^k share an edge, then (14) determines ϵ^{jk} up to a positive scalar multiple, since the shared edge has dimension (n-1) and $\delta^k-\delta^j$ has nonempty interior. This fact and Fenchel's inequality imply the following two lemmas.

<u>Lemma 7</u>: Suppose that Δ is a subgradient decomposition, and specifically that the convex function $V:M\to\mathfrak{R}$, the function $c:\Delta\to\mathfrak{R}$, and the 1-1 function $g:\Delta\to\mathfrak{R}^n$ satisfy (12) and (13). (That is if $\mu\in\text{int}(\delta)$, then $\nabla V(\mu)=g(\delta)$.) If

(15) $\epsilon^{jk}=g(\delta^k)-g(\delta^j)$,

and if E is the set of all of these vectors for δ^j and δ^k sharing a common edge, then (Δ,E) is a polyhedral directed graph.

<u>Lemma 8</u>: For every polyhedral graph Δ, at least one polyhedral directed graph (Δ,E) exists. If (Δ,E) and (Δ,F) are polyhedral directed graphs defined from Δ, having directed edges ϵ^{jk} and ξ^{jk} respectively, then there exist positive scalars α^{jk} such that

(16) $\epsilon^{jk}=\alpha^{jk}\xi^{jk}$.

---

generality in restricting our attention to polyhedral graphs.

Directed edges will now be defined for the polyhedral graph Δ of Example 2, and it will be shown that positive scalars α^{jk} cannot be found to satisfy (15) and (16) for any 1-1 function $g:\Delta\rightarrow\mathbb{R}^n$. By Lemma 7, then, Δ cannot be a subgradient decomposition.

One set F of directed edges satisfying (14) consists of: $\xi^{40}=(1,0)$, $\xi^{41}=(0,1)$, $\xi^{42}=(-1,0)$, $\xi^{43}=(0,-1)$, and $\xi^{01}=(-1,2)$, $\xi^{12}=(-2,-1)$, $\xi^{23}=(1,-2)$, $\xi^{30}=(2,1)$, and $\xi^{jk}=-\xi^{kj}$ for the remaining elements.

Suppose that a polyhedral directed graph could alternatively be induced on Δ by the subgradient vectors of a polyhedral convex function, as envisioned by Lemma 7. That is, suppose that $g:\Delta\rightarrow\mathbb{R}^n$ is 1-1, and that (15) and (16) define positive scalars α^{jk} (where the vectors ε^{jk} are as in Lemma 8). A contradiction will be derived. By (15),

(17) $0=\varepsilon^{40}+\varepsilon^{01}+\varepsilon^{14}$.

Applying (16) to (17) yields

(18) $0=\alpha^{40}\xi^{40}+\alpha^{01}\xi^{01}+\alpha^{14}\xi^{14}$.

The x-coordinate of (18) is

(19) $0=\alpha^{40}-\alpha^{01}$,

and the y-coordinate is

(20) $0=2\alpha^{01}-\alpha^{14}$.

Multiplying (19) by 2 and adding (20) yields

(21) $0=2\alpha^{40}-\alpha^{14}$.

By (15), $\varepsilon^{41}=-\varepsilon^{14}$. Also $\xi^{41}=-\xi^{14}$, so $\alpha^{41}=\alpha^{14}$ by (16). Applied to (20) and (21), this yields

(22) $\alpha^{41}=2\alpha^{40}$.

Similarly it can be shown that

(23) $\alpha^{40}=2\alpha^{43}$, $\alpha^{43}=2\alpha^{42}$, and $\alpha^{42}=2\alpha^{41}$.

But (22) and (23) imply that $\alpha^{40}=16\alpha^{40}$, which is only possible if $\alpha^{40}=0$. This would violate Lemma 8, though, so a contradiction has been reached. Thus Δ cannot be a subgradient decomposition, although it is a polyhedral graph.

In the context of a specific example, it has just been shown that a subgradient decomposition must satisfy a discrete version of an integrability condition. Integrability conditions are familiar from revealed-preference theory where, in the presence of subsidiary technical conditions, they are sufficient as well as necessary for behavior to be consistent with utility maximization. The present integrability condition will now be stated in general terms, and will be shown to be sufficient as well as necessary for a polyhedral graph to be consistent (in the sense of Theorem 1) with expected-utility maximization.

An elementary circuit in a directed graph is a path from some vertex, following edges, until the initial vertex has been reached again. A polyhedral directed graph is a directed graph, with the vertexes being the elements of Δ. The polyhedral graph Δ shown in Figure 2 was proved not to be a subgradient decomposition by proving that, in any polyhedral directed graph (Δ,E), the vector sum of the directed edges of some elementary circuit must be nonzero. In fact, the condition that E can be found for which every such vector sum equals zero is both necessary and sufficient for Δ to be a subgradient decomposition. This equivalence is stated as a theorem, after a formal definition of an elementary circuit is given.

A subset $\gamma \subseteq \Delta^2$ is an elementary circuit of the polyhedral graph Δ if there is a subset $D \subseteq \Delta$ such that

(24) $\gamma \subseteq \{(j,k) \mid j \in D \text{ and } k \in D \text{ and } \delta^j \text{ and } \delta^k \text{ share an edge}\}$,

(25) $\forall j \in D \ \exists! k \ (j,k) \in \gamma$,

(26) $\forall k \in D \ \exists! j \ (j,k) \in \gamma$, and

(27) No nonempty $\gamma' \subset \gamma$ satisfies (24), (25) and (26) with D replaced by any $D' \subseteq \Delta$.

Define Γ to be the set of elementary circuits of Δ.

<u>Theorem 3</u>: If Δ is a polyhedral graph, then Δ is a subgradient decomposition *if and only if* there exists a polyhedral directed graph (Δ, E) such that the sum of directed edges along every elementary circuit is zero, i. e.,

(28) $\forall \gamma \in \Gamma \ \Sigma_{(j,k) \in \gamma} \epsilon^{jk} = 0$.

Outline of the proof: Suppose that Δ is a subgradient decomposition, satisfying (10) with respect to the convex function V. Then E can be defined by

(29) $\epsilon^{jk} = \nabla V(\mu^k) - \nabla V(\mu^j)$,

where $\mu^j \in \delta^j$ and $\mu^k \in \delta^k$, whenever δ^j and δ^k share an edge. Conversely, suppose that (28) holds. A convex function V satisfying (10) will be constructed. At stage 1, define $V(\mu) = 0$ for $\mu \in \delta^1$ and $\nabla V(\mu) = 0$ for $\mu \in \text{int}(\delta^1)$. For $t > 1$, define $D_t = \{\delta \mid V \text{ and } \nabla V \text{ have been defined for } \delta \text{ by stage } t-1\}$. Suppose that $\Delta \backslash D_t \neq \emptyset$. It is easily shown that a polyhedral graph is connected, which implies that there are j and k such that δ^j and δ^k share an edge ω, $\delta^j \epsilon D_t$, and $\delta^k \notin D_t$. Let $\mu \epsilon \omega$, $\mu' \epsilon \mathrm{int}(\delta^k)$, and $\mu'' \epsilon \delta^k$. Then, at stage t, define ∇V on $\mathrm{int}(\delta^k)$ by (29) and define V on δ^k by

(30) $V(\mu'')=V(\mu)+\nabla V(\mu')(\mu''-\mu)$.

By (28), this recursive definition is consistent. The proof that V is convex is based on (14). Q.E.D.

The motivation for Theorem 3 is to determine whether the convex function required by (10) in Theorem 1 exists. By Theorems 2 and 3, this question can be answered by finding whether $\Delta=\{cl(\pi) | \pi \epsilon \Pi\}$ is a polyhedral graph, and, if so, whether a polyhedral directed graph (Δ,E) exists that satisfies (28). A reduction of this latter question to a system of linear weak inequalities is now obtained. In order to do this, first form a polyhedral directed graph (Δ,F) which is guaranteed by Lemma 8 to exist. (For instance, F may be constructed by taking unit vectors that satisfy (14) to be the directed edges ξ^{jk}.) Then, by Lemma 8, (28) can be satisfied if and only if there exist positive scalars α^{jk} (defined whenever δ^j and δ^k share an edge) such that

(31) $\forall \gamma \epsilon \Gamma \ \Sigma_{(j,k)\epsilon\gamma} \alpha^{jk} \xi^{jk}=0$.

By the linearity of the equations in (31), the existence of positive solutions is equivalent to the existence of α^{jk} that satisfy both (31) and also

(32) $\alpha^{jk} \geq 1$.

Thus, (31) and (32) define a system of linear weak inequalities, the consistency of which system is equivalent to Δ being a subgradient decomposition, and hence equivalent to there existing a utility function that rationalizes the partition Π from which Δ is derived.

6. Conclusion

Standard axiomatizations of expected-utility theory envision an agent with fixed probability assessments, who can be observed to choose actions from varying opportunity sets (for instance, pairs of lotteries) These axiomatizations also envision that the agent's preferences among these actions depend on the state of nature only through the state-dependent consequences of the actions, and that these consequences are clearly defined and observable. In this paper, we have suggested that this conception of the observable consequences of the theory may be an unsatisfactory basis for empirical testing of the theory. Rather, there may be many opportunities to observe choices from a fixed set of actions as probability assessments change. We have also suggested that the specification of utilities in terms of consequences may be quite difficult to implement empirically. On this view, a less restrictive specification of state-dependent utility needs to be studied.

It might be asked whether this new conception of observable consequences of expected-utility maximization places any testable restrictions on observation. The results of this paper show that there will be testable restrictions if the decision maker's probability assessments can be observed or estimated. One of these restrictions is that each action must be chosen for all probability measures in the interior of some convex subset of the probability simplex. The example constructed in section 5 (and depicted in Figure 1) shows that this convexity restriction does not exhaust the empirical content of the theory. Theorem 3 fully characterizes this empirical content in terms of a discrete version of an integrability condition, familiar from the theory of choice under certainty.

There are at least two ways in which the decision maker's probability assessments might be treated as observable or estimable. One would be to elicit them directly, perhaps using a "proper scoring rule" that would provide incentives for truthful revelation. The other would be to impose a rational-expectations assumption that the decision maker knows the true statistical distribution from which observations are drawn. Consider, for instance, the example of medical diagnosis that was discussed in the introduction. The medical history and examination results that constitute the doctor's evidence about a patient can be viewed as a discrete-valued measurement. This measurement partitions the set of patients into finitely many equivalence classes. That is, all patients within a class have identical medical histories and examination results. Within each equivalence class, there is a sample distribution of eventual, accurate diagnoses.[16] These diagnosis outcomes can be treated as

---

[16]Edwards (1972, pp. 139-140) quoted by Shortliffe and Buchanan (1984, p.236) has described some of the practical difficulties of gathering and interpreting such a data set. Nevertheless, particularly because the use of computers has dramatically improved the quality of medical data available, the

an independent sample drawn from the doctor's posterior probability distribution conditioned on the evidence. The likelihood function regarding the beliefs of the doctor can then be formed in a standard way. This rational-expectations approach to estimation has two particular virtues when the *decision maker* is an expert in the decision domain being studied. First, such an expert will have substantial experience, and standard results of decision theory imply that conditioning on such experience should closely approximate the distribution of the evidence. Second, regardless of what are the subjective beliefs of the decision maker, whether an expert optimizes with respect to the empirical distribution of a large sample of decision situations is of *independent interest*. This objective assessment, rather than the expert's own subjective assessment of success, might even be what the expert's clients are most eager to know.

Finally, it might be asked whether the characterization of expected-utility maximization provided in Theorem 3 can potentially be of direct use in the kind of empirical study that is suggested here. This question arises because any data set will be finite, so its consistency with expected-utility theory can be determined directly by seeing whether finitely many instances of inequality (2) are a consistent system. The answer to the question has to do with computational complexity. Let α be the number of actions available to the decision maker. Then the number of linear inequalities that must be solved to

---

use of such data may well be feasible.

implement the direct approach is $(\alpha-1)$ times the cardinality of the data partition that was defined in the preceding paragraph. Theorem 3 concerns a polyhedral graph that has at most $\alpha(\alpha-1)/2$ edges, and the number of linear inequalities that have to be solved is roughly equal to the number of states of nature times the number of elementary cycles of this graph. If the sets of states of nature and of actions are small, but the data partition is fine, it might be possible to apply the geometric insights of Theorem 3 to minimize the amount of computation that is required. In particular, only comparisons between pairs of actions corresponding to adjacent vertexes of the polyhedral graph, and comparisons made at partition elements having sample distributions close to the boundaries of those vertexes, may be necessary. *If so, then Theorem 3 may have considerable practical importance.*

Figure 1

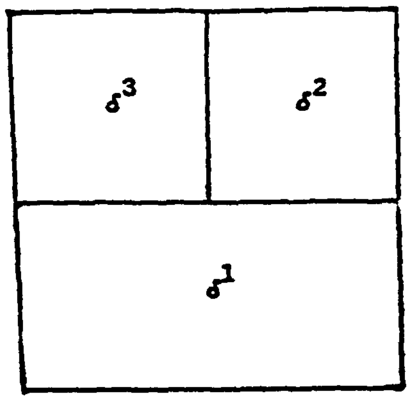

Figure 2

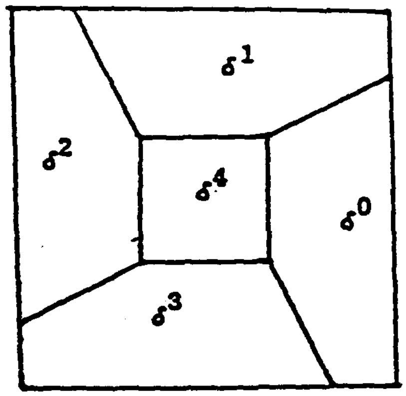

Figure 4

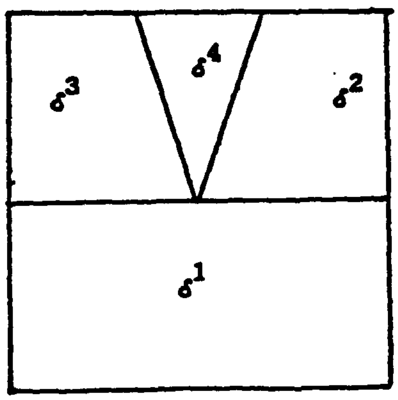

Figure 3

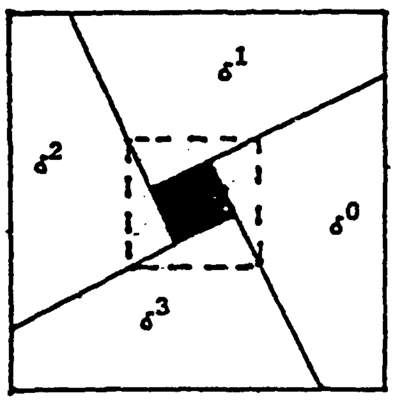

REFERENCES

Allais, M., "Allais' rejoinder: Theory and empirical evidence," in Allais and Hagen, eds., Expected Utility Hypothesis and the Allais Paradox, Dordrecht; Boston: D. Reidel, 1979.

Attanasio, O. P., and G. Weber, "Intertemporal substitution, risk aversion and the Euler equation for consumption: Evidence from aggregate and average cohort data," University College, London, 1987.

Border, K., "Stochastic dominance and choice of lotteries," California Institute of Technology, 1987.

Chew, S. H., L. Epstein, and I. Zilcha, "A correspondence principle between expected utility and smooth utility," Tel-Aviv University, 1986.

Edwards, W., "N=1: Diagnosis in unique cases," in Jacquez, ed., Computer Diagnosis and Diagnostic Methods, 1972.

Epstein, L., and S. Zin, "Substitution, risk aversion and the temporal behavior of asset returns," University of Toronto, 1987.

Green, J. R., L. Lau, and H. Polemarchakis, "Identification of the von Neumann-Morgenstern utility function from asset demands," in General Equilibrium, Growth and Trade, J. R. Green and J. Scheinkman, eds., New York: Academic Press, 1978.

Green, R. C., and S. Srivastava, "Expected utility maximization and demand behavior," J.E.T., 1986.

Grether, D. M., and C. R. Plott, "Economic theory of choice and the preference-reversal phenomenon," A.E.R., 1979.

Karni, E., D. Schmeidler, and K. Vind, "On state-dependent preferences and subjective probabilities," Econometrica, 1983.

Machina, M. J., "Expected utility analysis without the independence axiom," Econometrica, 1982.

Miller. R., "Job matching and occupational choice," J.P.E., 1984.

Osband, K., and S. Reichelstein, "Information-eliciting compensation schemes," Journal of Public Economics, 1985.

Osband, K., "Eliciting forecasts: a general characterization of scoring rules," manuscript, 1986.

Pakes, A., "Patents as options: some estimates of the value of holding European patent stocks," Econometrica, 1986.

Richter, M. K., "Duality and rationality," JET, 1979.

Rockafellar, R. T., Convex Analysis, Princeton: Princeton University Press, 1970.

Rust, J., "Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher," Econometrica, 1987.

Savage, L. J., "Elicitation of personal probabilities and expectations," JASA, 1971.

Savage, L. J., The Foundations of Statistics, New York: Wiley, 1954.

Shortliffe, E. H., and B. G. Buchanan, "A model of inexact reasoning in medicine," in Buchanan and Shortliffe, eds., Rule-Based Expert Systems, 1984.

Varian, H., "Nonparametric tests of models of investor behavior," J.F.Q.A., 1983.

Wolpin, K. I., "An estimable dynamic stochastic model of fertility and child mortality," J.P.E., 1984.